

COMPUTING ENVIRONMENTS FOR DATA ANALYSIS

Part 2: Hardware

John Alan McDonald


and

Jan Pedersen

Technical Report No. 10

April 1984

**Laboratory for
Computational
Statistics**



**Department of Statistics
Stanford University**

This document and the material and data contained therein, was developed under sponsorship of the United States Government. Neither the United States nor the Department of Energy, nor the Office of Naval Research, nor the U.S. Army Research Office, nor the Leland Stanford Junior University, nor their employees, nor their respective contractors, subcontractors, or their employees, makes any warranty, express or implied, or assumes any liability or responsibility for accuracy, completeness or usefulness of any information, apparatus, product or process disclosed, or represents that its use will not infringe privately-owned rights. Mention of any product, its manufacturer, or suppliers shall not, nor is it intended to, imply approval, disapproval, or fitness for any particular use. A royalty-free, nonexclusive right to use and disseminate same for any purpose whatsoever, is expressly reserved to the United States and the University.

COMPUTING ENVIRONMENTS FOR DATA ANALYSIS*

Part 2: Hardware

John Alan McDonald

and

Jan Pedersen

Department of Statistics

and

Stanford Linear Accelerator Center

Stanford, California

LCS Technical Report No. 10

April 1984

ABSTRACT

This is the second in a series of papers on aspects of modern computing environments that are relevant to statistical data analysis. We argue that a network of graphics workstations is far superior to conventional batch or time-sharing computers as an environment for interactive data analysis. The second paper in the series discusses the hardware of graphics workstations in detail.

* Work supported by a NSF Mathematical Sciences Postdoctoral Research Fellowship, Office of Naval Research contract N00014-83-K-0472 and grant N00014-83-G-0121, and U.S. Army Research Office under contract DAAG29-82-K-0056

Computing Environments for Data Analysis Part 2: Hardware

John Alan McDonald

and

Jan Pedersen

Dept. of Statistics

and

Computation Research Group

Stanford Linear Accelerator Center

Stanford University

April 1984

Abstract

This is the second in a series of papers on aspects of modern computing environments that are relevant to statistical data analysis. We argue that a network of graphics workstations is a far superior to conventional batch or time-sharing computers as an environment for interactive data analysis. The second paper in the series discusses the hardware of graphics workstations in detail.

Support

This research was supported by a NSF Mathematical Sciences Postdoctoral Research Fellowship, Office of Naval Research contract N00014-83-K-0472, Office of Naval Research grant N00014-83-G-0121, and U.S. Army Research Office under contract DAAG29-82-K-0056.

Keywords: *Data Analysis, Interactive Graphics, Workstations, Programming Environments*

Chapter 1

Introduction

The purpose of this paper is to list important features to consider when evaluating and comparing graphics workstations to be used for data analysis research.

The workstations described here cannot be used for data analysis "off the shelf"; they are at present useful primarily for research. We expect a minimum of 5-10 years of development before one will be able to buy both a workstation and statistical software to use on it.

The computer graphics market changes rapidly; new, better, cheaper machines are announced every few months. In the course of this paper, we will mention features of six specific machines—chosen somewhat arbitrarily out of the hundreds that are available—as concrete examples. Four of these, Apollo, Chromatics, Iris, and Sun, are representatives of a large class of graphics workstations based on versions of the Motorola 68000 microprocessor. This type of workstation typically costs \$40,000 to \$60,000 in configurations appropriate for statistical applications. The fifth machine, the Ridge 32, has a *reduced-instruction-set* central processor that is a bit more expensive and considerably more powerful than a 68000. The architecture of the sixth machine, the Symbolics 3600, is specially designed for efficient execution of programs written in Lisp; it is both more powerful and more expensive (both by a factor of 2-3) than the 68000 based machines.

We emphasize that this is by no means an exhaustive list; these are machines which have been, are being, or will be used for research in data analysis systems. Apollo is used by the Statistics Dept. at Harvard; Chromatics by the Statistics Depts. at Stanford and Berkeley; the Iris will be used by Statistics Dept. at Stanford; the Ridge32 will be used by Statistics Depts. at Stanford, Berkeley, and the U. of Washington; Sun hardware was the basis of the **Orion I** workstation at Stanford; the Symbolics 3600 is used by the Statistics Depts. at Stanford and the U. of Washington, and by groups at Bell Labs.

A reference for micro-computer architecture is Kraft and Toy [1979]. Foley and VanDam [1982] discuss the architecture of microcomputer graphics systems.

Chapter 2

The Central Processor

2.1 Single-chip Microprocessors

One or two dozen companies make graphics workstations based on the Motorola 68000 microprocessor [Motorola, 1982]. Examples are Apollo, Chromatics, Iris, and Sun. In the future, other 32 bit microprocessors will be important competitors (for example, the National Semiconductor NS16032 and NS32032, the Hewlett-Packard HP-9000, Bell Labs' Bellmac 32, Zilog Z80,000), but there are few available workstations based on these processors (as of April 1984).

The 68000 is sometimes decribed as a 16/32 bit machine; internally it is a 32-bit machine; it communicates with the outside world with 16-bit data words and 23-bit address words. Motorola has recently released a version called the 68010, which supports virtual memory (the plain 68000 does not). Iris and Sun, and possibly also the Apollo will use the 68010; the Chromatics uses the 68000. A fully 32-bit version of the 68000 (the 68020) is promised for the future.

2.2 Bit-sliced Processors

Bit-sliced processors are an alternative to one-chip microprocessors like the 68000. Bit-sliced processors are made by combining chips which perform standard operations on a *slice* of a machine word. Slices are commonly four bits wide; so, for example, eight 4-bit adders would be combined to make a 32-bit adder, which would be combined with other 32-bit parts to make a complete processor. Bit-sliced processors usually execute a *micro-coded* instruction set. That is, each assembly language instruction calls a small, fast program written in a simpler *microcode* language, which is executed directly by the bit-slice hardware.

Bit-slice components are often used for special purpose processors; the flexible hardware and the micro-coded instruction set make it possible to design and construct

a new processor without the capital investment needed to produce a new microprocessor. Bit-slice processors are usually faster (and more expensive) than general-purpose, single-chip microprocessors because they are designed to be particularly efficient at special tasks and also because they typically use faster transistor technology.

2.2.1 Reduced Instruction Set Machines

The *Ridge32* is a graphics workstation with a bit-sliced central processor that is roughly equivalent to a VAX 11/780, or 2-3 68000's, for a price that is no more than that of typical 68000 based workstations. The low price and high power of the Ridge is due, in part, to the fact that it is a *reduced instruction set computer (RISC)*. This means that it has a relatively simple instruction set, which is executed directly by the hardware, rather than going through *micro-code translation* first.

2.2.2 Lisp Machines

The *Symbolics 3600* [Symbolics, 1983; Bawden, et al., 1979] is another graphics workstation with a bit-sliced central processor. The 3600 is a *LISP* machine, that is, its architecture is designed for fast execution of programs written in LISP. This has, obviously, implications for the 3600's software properties, which will be discussed below. The 3600 is equivalent to 1 or 2 VAX 11/780's, or 2-5 68000's and it is about 3 times the price of a typical 68000-based workstation.

Other Lisp machines are:

Lisp Machine, Inc.'s *Lambda* is very similar to the 3600, in power and price, with the addition of an auxilliary 68000 processor running Unix.

Xerox has its *D- or 1100-series*, which ranges from the *Dandelion*, comparable to a 68000 based machine, to the *Dorado*, [Dorado, 1981] which is similar to a 3600.

2.3 Clock speed and wait states

The *clock speed* of a 68000 is its basic cycle time and, therefore, determines how fast it executes instructions. The clock speed is a property of the computer in which the 68000 chip is placed. It is also a property of the chip, in the sense that not all chips will run successfully at high speeds. Typical speeds range from 4 to 20 megahertz. A 16mhz machine is, not surprisingly, about 4 times as fast as a 4mhz one. Every 16mhz machine is not exactly 4 times as fast as every 4mhz one because the speed of memory and the *bus* that connects the 68000 to the memory can make a considerable difference to the effective speed of the processor. Bus and memory speed are reflected in the *number of wait states* in a given machine, which are processor cycles lost while the processor waits for the bus. As a rough guide to the power of 68000 machines, the Sun

workstation contains a 10 mhz 68010 with no wait states and is roughly comparable to a VAX 11/750 (for non-floating computations).

Chapter 3

Auxilliary Processors and Memory

3.1 Bus Architecture

The *bus* is used by the central processor to communicate with memory and other I/O devices, in particular, the graphics display. A common standard bus used in 68000 based machines is the MULTIBUS [Boberg,1980], used in Sun and Iris. It is an advantage to have a machine with a standard bus architecture, because it is then easier to add peripherals, such as floating point boards, array processors, printers, or additional input devices.

Some machines (Sun and Iris, for example) are designed so that the RAM is either on the same board as the 68000 or is connected with a separate, private bus. This allows the 68000 to access memory without using the main (MULTIBUS) bus, avoiding delays that result when the bus must be shared with auxilliary processors, such as a floating point board or array processor.

Another reason that some workstations add a second, private bus is that the MULTIBUS is not fast enough for a 68000 run at high clock speeds. In the future, workstations will be provided with faster 32 bit buses, to keep up with 32 bit processors run at high clock speeds.

Many workstations increase their computational power by adding auxilliary processors which may be more efficient at specialized tasks and execute in parallel with the central processor.

3.2 Scalar floating point

The 68000, like most other microprocessors, does not include floating point operations in its instruction set. A machine with only a 68000 therefore executes floating point operations as small programs, slowly. Efficient software floating point operations take 50—100 microseconds on a 10mhz 68000 with no wait states. The same

operations require closer to one microsecond on a VAX. One solution is to add an auxillary floating point processor to a 68000-based machine. A floating point board is included in the Apollo and is available (from Sky Computers, Inc., for example) for the Chromatics, Iris, and Sun.

The Ridge32 and the 3600 bit-sliced central processors do not have floating point operations as primitives. However, software floating point is efficient enough to give about half the speed of a VAX 11/750 for the Symbolics and speed about equal to the 11/750 on the Ridge. Both Ridge and Symbolics plan to add hardware floating point processors in the future.

3.3 Vector Floating Point

Many graphical and statistical computations can be efficiently done in parallel. So it is useful to have a workstation with an *array processor*. The Iris workstation includes an array processor based on a chip called the *Geometry Engine* [Clark, 1982], which is specially designed for graphical computations and should be able to perform about 10 million floating operations per second—the equivalent of 10-20 VAX 11/780's. General purpose array processors (from Sky Computers, Inc. for example) can be added to Chromatics, Sun, and Iris.

The other machines have non-standard bus architectures, which makes it difficult to add array processors that can communicate quickly with the cpu.

3.4 I/O Processors

Some workstations contain special processors for input and output, which free the central processor to do other work. I/O processors may, for example, control the keyboard, handle data transfers between disk and memory (DMA's), buffer communications with a network, or speed communication with the graphics device (see below).

3.5 Random Access Memory

Graphics workstations are usually provided with from 1/2 to 2 megabytes of *RAM* (*random access memory*). The speed of the memory places a constraint on the effective speed of computation. Slow memory will require the processor to have some number of wait states (lost cycles, mentioned above).

3.6 Memory Caches

Some workstations gain additional processing speed by including *high speed caches* [Clark, Lampson, and Pier, 1981; Smith, 1983]: *instruction pre-fetch caches* and/or *data stack caches*. A cache is a limited amount of high speed memory; by keeping the next instructions to be executed and/or currently relevant data in a cache a computer can increase its effective rate of computation. The Apollo has an instruction cache; the 3600 has both instruction and data caches.

3.7 Virtual Memory

Virtual memory allows a computer to run programs and manipulate data sets too large to fit in its physical memory. It is a property of the computer's operating system as well as its hardware. Some processors cannot run virtual memory operating systems (because they cannot recover from page faults). The original 68000 cannot run a virtual memory operating system. The more recent 68010 can. The Chromatics uses the 68000 and does not have virtual memory; Iris and Sun use the 68010. Older versions of Apollo use two 68000's to support virtual memory in a somewhat clumsy way; future Apollos should use the 68010. The Ridge and 3600 cpus support virtual memory.

3.8 Disk Storage

Graphics workstations usually include a disk for long term data storage; most manufacturers offer a range disk sizes, typically from 10 to 500 megabytes. Workstations use *Winchester* disk technology, which permits disks to be physically much smaller and less expensive than traditional disk technology. 50 megabytes is adequate for statistical applications; 500 megabytes is not an unreasonable size. A 500 megabyte disk will cost two or three times a 50 megabyte one—from \$15,000 to \$30,000.

Chapter 4

Graphics Device

4.1 Bitmap Graphics

All of the graphics workstations that we are considering here use a *bitmap graphics display* which is also called a *raster graphics display* or a *frame buffer*. For a fuller discussion of bitmap graphics devices see: a survey paper aimed at statisticians by Beatty [1983] and texts by Foley and VanDam [1982] and Newman and Sproull [1979].

4.1.1 Resolution

In a bitmap graphics device, the picture is made up of an array (a raster) of dots, called *pixels*. The *resolution* of the device is the number of pixels on the screen. Common resolutions are either approximately 512x512 or 1024x1024. For many kinds of statistical graphics 512x512 is enough. A 1024x1024 bitmap is needed for realistic images of solid objects, scatterplots in which small symbols (glyphs) are used to code additional variables, or if the screen is used to display several plots at once.

4.1.2 Depth

The color of each pixel is determined by the part of the *refresh memory* associated with it. The number of bits of refresh memory associated with each pixel is the *depth* of the bitmap. This is also referred to as the *number of bitplanes* in the graphics device. Typical depths are 1, 4, 8, 16, 24, or 32, with 1 and 8 being the most common.

The depth of the bitmap determines the number of different colors that can be displayed simultaneously. A device with a single bitplane can display only two colors (e.g. black and white) depending on whether the single bit of refresh memory corresponding to a pixel is 0 or 1. A device with eight planes can display 256 colors at once.

Obviously, the more planes a bitmap has the better. Adding planes increases the price of a system; for a 1024x1024 bitmap, each additional eight planes requires an additional megabyte of refresh memory. The refresh memory must be fast enough to permit its entire contents to be read 30 times per second (60 times a second for non-interlaced monitors). It is therefore more expensive than ordinary RAM; refresh memory prices are \$7,000—\$20,000 per megabyte as compared to \$2,000—\$7,000 per megabyte of RAM.

For most statistical applications, eight planes is enough. Bitmaps with more depth are useful for drawing realistic solid objects and for image processing applications. Extra bitplanes are also useful for *double buffering*, drawing multiple independent images, and *color map animation*.

4.2 Refresh Cycle

A bitmap graphics device maintains a picture on the screen by going through the *refresh cycle* fast enough to avoid flicker (typically either 30 or 60 times per second).

In a bitmap device, the electron beam scans over the screen in a regular fashion, modulating the intensity of the beam as it goes to determine the brightness and color of each pixel. A raster line is usually drawn as the beam scans from left to right. With the intensity set to zero, the beam then moves back quickly, from right to left, to the start of the next raster line (the *horizontal blanking or retrace* part of the refresh cycle). When the beam reaches the bottom of the screen, it quickly move back to the top, with the intensity at zero (*vertical retrace*). Some devices scan every other raster line in one pass from top to bottom, getting the missed lines in a second pass (see section on *interlacing* below).

When the electron beam strikes a point on the screen, the phosphors emit light with a brightness that depends on the intensity of the beam. In color systems, there are small dots of red, green, and blue phosphors, whose emissions mix to produce pixels that are perceived as having arbitrary colors. The brightness of the phosphors decays rapidly after the electron beam leaves. To produce a stable picture, the electron beam must return before the decay in brightness is noticeable.

The picture on the screen is determined by the contents of the refresh memory. With a 1024x1024x8 bitmap and a 30hz monitor, the refresh memory must be read at about 32 megabytes per second (to allow time for horizontal and vertical retrace). The 8 bits per pixel is decoded into, typically, 8 bits each of red, green, and blue intensity by a color map (next section). *Digital-to-analog converters* translate the color intensities to voltages that modulate the intensity of the electron beam.

4.3 Color Maps

Color maps (also called *color look up tables*) are used to make bitmap displays more flexible.

A primitive bitmap display with, say, eight planes might assign a fixed color to each of the possible pixel values (0-255). However, it is often convenient if the color associated with each pixel value can be changed. This is done using a color map. Typical color maps are *8-bits-in 24-bits-out*; they can be thought of as three arrays, one each for red, green, and blue; each array has 256 entries, indexed from 0 to 255; the pixel value is used as an index into these arrays to determine the relative intensities (from 0-255) of red, green, and blue.

4.3.1 True Color

True color is an alternative to a color map. True color may be used, for example, on a system with 24 bitplanes. A 24-bit-in 24-bit-out look up table is impractical, because it would require $3 * 2^{24}$ bytes of very high speed memory. Instead, the 24 bit (3 byte) pixel values are used to hold 8 bits of intensity for red, green, and blue. True color is available for the Iris and the 3600.

4.3.2 Double Buffering

Graphics systems often use extra (more than 8) bitplanes for *double buffering*. A 16 plane frame buffer can hold two independent 8 plane pictures; in some machines (Chromatics), it is possible to quickly change which 8 planes are displayed on the screen. Double buffering is useful in motion graphics; a new picture is drawn into the 8 planes that are invisible, the buffers are swapped so that the new picture is visible and the old picture is invisible, and the invisible 8 planes are erased. This eliminates distracting *beating* or *aliasing* effects that arise when the drawing-erasing cycle is visible and interferes with the refresh cycle.

4.3.3 Synchronization of Refresh and Color Map Changes

A slightly subtle consideration is whether changes to the color map are synchronized with the refresh cycle. The color map should only be changed when the screen is in the blanked part of the refresh cycle. Otherwise there will be contention between the refresh and the cpu for access to the color map, which can cause disturbing effects on the display.

4.4 Graphics Speed

Many statistical applications require motion graphics, so the time required to modify the bitmap—the time required to erase an old picture and draw a new one—is important. The time required to change a picture should simply be the time required to change a single pixel times the number of pixels that have to be changed. However, the time to change a pixel can vary greatly depending on how it is changed. Common modes for writing to a bitmap are: single pixel change, vector drawing, rectangular area fill, polygonal area fill, and an operation on rectangular blocks of pixels called *RasterOp* or *BitBlit* [Bechtolsheim and Baskett, 1980]. Typical speeds are 0.1 to 2 million pixels per second in single pixel write, 0.1 to 16 million pixels per second in vector drawing, and up to 200 hundred million pixels per second in area fill (a special area fill mode on the 3600). For comparison, to redraw an entire screen in real-time requires drawing 1 million pixels 30 times a second, or 30 million pixels per second. To draw a rotating 3-dimensional scatterplot containing 1000 points, in which each point is represented by a symbol composed of, say, 5 vectors 10 pixels long, would require a vector drawing speed of at least 3 million pixels per second.

4.4.1 Graphics Processors

The wide range in drawing speeds is in part due to the fact that some devices have special graphics processors to speed up some drawing tasks.

For example, to draw a vector, it must be *rasterized*, that is, some processor must decide which pixels have to be written to connect the vector's two endpoints. On the Sun, for example, this computation is done by the 68000, which makes vector drawing slow. The Chromatics, in contrast, contains a special vector processor, which rasterizes the vector and modifies the necessary pixels, freeing the 68000 for other computation.

The Iris is an extreme example of a machine with auxilliary graphics processing. The Iris includes the VLSI *Geometry System* [Clark, 1982] a powerful, general purpose graphics processor which will generate linear vectors, quadratic and cubic curves, all conic sections, rotate, scale, clip, do the perspective computation, etc.

4.4.2 Communication between Central Processor and Bitmap

The speed of picture drawing is limited by the speed and intimacy of communication between the cpu and the bitmap. On some systems (Chromatics and 3600, for example) the refresh memory is on the cpu bus and is in the cpu's address space; in other words, the pixels in the bitmap can be accessed by the cpu as though they were bytes of ordinary RAM. This makes communication between the cpu and the bitmap fast and flexible. In a less desirable alternative, the cpu talks to the bitmap

by sending graphics commands to a special interface (this was the case on **Orion I**), which is relatively slow and awkward.

4.5 Monitor

A 19 inch color monitor is standard on graphics workstations. A monitor may run *30 hz interlaced*, which means that all the even scan lines are drawn in 1/60th of a second and the odd scan lines are drawn in the next 1/60th of a second, so that the entire screen is refreshed 30 times a second. A better but more expensive alternative is a *60hz non-interlaced* monitor, which draws all the scan lines on the screen, in one pass, 60 times a second.

Another consideration in the choice of a monitor is the *persistance* of the phosphors. Each point on the screen is hit by the electron beam 30 or 60 times a second. The brightness of a point decays exponentially after the electron beam leaves it. The phosphors must be persistent enough so that the screen does not flicker noticeably in the 1/30 of a second between refreshes. On the other hand, if the phosphors are too persistent, moving pictures (rotating scatterplots) will produce distracting streaks.

60hz non-interlaced monitors are less likely to have objectionable *flicker*, making shorter persistence phosphors practical. With 30hz interlaced monitors, flicker may not be a problem, depending on room lighting and a variety of other factors. In some machines (Chromatics), the refresh is synchronized with the 60hz cycle in the wall outlet's alternating current. Thus the monitor is synchronized with the 60hz oscillation in the brightness of florescent room lighting, which helps minimize flicker.

4.6 Miscellaneous Features

Graphics devices come with a variety of additional features which are not critical to statistical applications, but are often useful nonetheless. Among these are:

- **Zoom**—magnify the bitmap so that a fraction (1/4, 1/9, 1/16 etc.) of it fills the screen.
- **Pan**—translate the bitmap after zooming, so that a different portion is visible.
- **Blink planes**—are used to make pixels blink; this is often done by alternately masking and not masking the blink bit, which causes the pixel value to alternate between two address in the color map.
- **Overlay planes**—are used to write text non-destructively over the bitmap.
- **Hardware cursor**—draws and moves a small pointing symbol or *cursor* non-destructively over the bitmap.

Chapter 5

Input Devices

In the preceding sections we have discussed the display, the workstation's output device. One advantage of graphical workstations is use of *graphical input devices*, which can make using a computer a natural activity, in the same way that driving a car is natural.

Foley and vanDam [1982] define five *logical input devices*: the *locator*, which indicates position and/or orientation, the *pick*, which selects an object displayed on the screen, the *valuator*, which chooses a real number, the *button(s)*, which select from a finite set of alternatives, and the *keyboard*, which inputs a character string.

Graphics workstations are usually provided with keyboards and some number of other graphical input devices, such as: joystick (Chromatics), mouse (Iris, Sun, 3600), and touchpad (Apollo). Each of these physical input devices, including the keyboard, can be used to implement any of the five logical input devices, but the implementation will be much more natural in some cases than others.

Chapter 6

Networking

Because graphics workstations are single-user, stand-alone computers, it is important for them to be able to communicate quickly and easily with other computers—other workstations or mainframes. Computers communicate better if they are linked in a *network* [Green, 1982; Tannenbaum, 1981a, 1981b]. A network can be used simply to transfer files between machines, or it can be used in a more sophisticated fashion to distribute computing tasks among machines in the network.

A common standard is Xerox's Ethernet, which is supported by Iris, Sun, and the 3600. Apollo supports its own, idiosyncratic network. The Chromatics has no networking support.

An example of more sophisticated use of a network is the Sun's option for diskless workstations. These diskless stations use the Ethernet to simulate a virtual disk which is physically part of a large disk maintained by a special *fileserver* station. A network of diskless stations provides more computing power for less initial expense and less maintainance, at the cost of somewhat slower disk access and potential bottlenecks when many stations need to use the central disk at the same time.

Chapter 7

Benchmarks

Many factors, software as well as hardware, determine the effective speed of computation. To get a more valid comparison of machines, it is useful to run one or several *benchmark* programs. A simple benchmark used to test floating point speed is to sum up the harmonic series (suggested by Peter Huber).

To sum 100,000 terms of the harmonic series (in single precision) takes approximately:

- 1 second on a Ridge32 (with software floating point).
- 1 second on a quiet (single user) VAX 11/750 (with floating point accelerator).
- 2 seconds on a Symbolics 3600 (without floating point accelerator and without high speed instruction pre-fetch cache). On a Symbolics 3600 with floating point accelerator and high speed cache it would presumably take much less than one second.
- 10 seconds on an Apollo (with hardware floating point and high speed cache).
- 35 seconds on a Sun (software floating point).
- 80 seconds on a Chromatics (software floating point).

Chapter 8

Sources

The best way to survey the computer graphics market is to attend annual meetings, such as the COMDEX convention, the NCC meeting, the National Computer Graphics Association (NGCA) meeting, usually in June, and the annual ACM SIGGRAPH meeting, usually in July.

The next best way is to review journals that carry advertisements and announce new computing and graphics products, such as **Electronics** magazine, **IEEE computer**, **IEEE Computer Graphics and Applications**, **Computer Graphics World**, etc.

The manufacturers of the workstations mentioned in this paper are:

Symbolics, Inc. (**3600**); 845 Page Mill Road; Palo Alto, California 94304; 415-494-8081.

Apollo Computer, Inc; Chelmsford, Massachusetts 01824; (617) 256-6600.

Chromatics, Inc. (**CGC-7900**); 2558 Mountain Industrial Boulevard; Tucker, Georgia 30084; (404) 493-7000.

Silicon Graphics, Inc. (**Iris**); 630 Clyde Court; Mountain View, California 94043; (415) 960-1980.

Sun Microsystems, Inc.; 2550 Garcia Ave.; Mountain View, California 94043; (415) 960-1330.

Ridge Computers (**Ridge 32**); 586 Weddell Drive; Sunnyvale, California 94089; (408) 745-0400.

Chapter 9

References

- Bawden, A., Greenblatt, R., Holloway, J., Knight, T., Moon, D., and Weinreb, D. (1979)**
The LISP Machine
in *Artificial Intelligence: An MIT Perspective, vol.II*
Winston, P.H. and Brown, R.H., eds. Cambridge, Mass., MIT Press, 1979
- Beatty, J.C. (1983)**
Raster Graphics and Color, *The American Statistician*, Vol. 37, Num. 1, Feb. 1983, pp 60-75
- Boberg, R.W. (1980)**
Proposed Microcomputer System 796 Bus Standard
Computer, Vol 13, Num. 10, Oct. 1980, pp 89-106
- Clark, J.H. (1982)**
The Geometry Engine: A VLSI Geometry System for Graphics
Proceedings 1982 SIGGRAPH Conference, published as: *Computer Graphics* 16 (3). pp 127-133
- Dorado (1981)**
The Dorado: A High Performance Personal Computer, Three Papers
Xerox PARC Report CSL-81-1. Palo Alto, CA
- Foley, J.D. and Van Dam, A. (1982)**
Fundamentals of Interactive Computer Graphics
Reading, Mass., Addison Wesley, 1982
- Green, P.E., ed. (1982)**
Computer Network Architectures and Protocols
New York, Plenum Press, 1982
- Kraft, G.D. and Toy, W.N. (1979)**
Mini/Microcomputer Hardware Design
Englewood Cliffs, N.J., Prentice Hall, 1979
- Motorola, Inc. (1982)**

MC68000, 16-Bit Microprocessor, User's Manual, Third Ed.
Englewood Cliffs, N.J., Prentice-Hall, 1982
Newman, W.M. and Sproull, R.F. (1979)
Principles of Interactive Computer Graphics, 2nd ed
New York, McGraw Hill, 1979
Symbolics Inc. (1983)
3600 Technical Summary
Symbolics, Inc. 5 Cambridge Center, Cambridge, Mass. 02142.
Tannenbaum, A.S. (1981a)
Computer Networks
Englewood Cliffs, N.J., Prentice Hall, 1981
Tannenbaum, A.S. (1981b)
Network Protocols
ACM Computing Surveys, Vol. 13, Num. 4, Oct. 1980, pp 89-106